

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	4	(multi\$1dimension\$2 or dimension\$2 or (two adj dimension\$2)) near2 queue same (priorit\$3 same time) and ((@ad<"20010618") or (@prad<"20010618") or (@rlad<"20010618"))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/07 12:45
L3	2	718/102-103.ccls. and (multi\$1dimension\$2 or dimension\$2 or (two adj dimension\$2)) near4 queue and ((@ad<"20010618") or (@prad<"20010618") or (@rlad<"20010618"))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/07 12:46
L4	2	718/102-103.ccls. and (multi\$1dimension\$2 or dimension\$2 or (two adj dimension\$2)) near5 queue and ((@ad<"20010618") or (@prad<"20010618") or (@rlad<"20010618"))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/07 12:46
L5	128	dimension\$2 near2 queue	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/07 12:53
L6	134	(multi\$1dimension\$2 or dimension\$2 or (two adj dimension\$2)) near2 queue	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/07 12:53
L7	105	(multi\$1dimension\$2 or dimension\$2 or (two adj dimension\$2)) near2 queue and ((@ad<"20010618") or (@prad<"20010618") or (@rlad<"20010618"))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/07 12:53
L8	61	(multi\$1dimension\$2 or dimension\$2 or (two adj dimension\$2)) near2 queue and priorit\$3 and ((@ad<"20010618") or (@prad<"20010618") or (@rlad<"20010618"))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/07 12:53
L9	60	(multi\$1dimension\$2 or dimension\$2 or (two adj dimension\$2)) near2 queue and priorit\$3 and time and ((@ad<"20010618") or (@prad<"20010618") or (@rlad<"20010618"))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/07 12:53

L10	3	"6609161".pn.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/11/07 12:53
L11	7	("5564023" "5625800" "5627745" "5797034" "5938747" "6006292" "6049526").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2005/11/07 12:53
L12	1	"20020194249"	US-PGPUB; USPAT	OR	OFF	2005/11/07 13:22
L13	1	"6609161".pn.	US-PGPUB; USPAT	OR	OFF	2005/11/07 13:25



USPTO

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
 The ACM Digital Library The Guide

two dimensional sleep run queue

THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used [two dimensional sleep run queue](#)

Found 81,969 of 166,357

Sort results
by

relevance

 [Save results to a Binder](#)Display
results

expanded form

[Search Tips](#) [Open results in a new window](#)[Try an Advanced Search](#)[Try this search in The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

1 A SMART scheduler for multimedia applications

Jason Nieh, Monica S. Lam

May 2003 **ACM Transactions on Computer Systems (TOCS)**, Volume 21 Issue 2**Publisher:** ACM PressFull text available: [pdf\(570.87 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Real-time applications such as multimedia audio and video are increasingly populating the workstation desktop. To support the execution of these applications in conjunction with traditional non-real-time applications, we have created SMART, a Scheduler for Multimedia And Real-Time applications. SMART supports applications with time constraints, and provides dynamic feedback to applications to allow them to adapt to the current load. In addition, the support for real-time applications is integrat ...

Keywords: Scheduling, multimedia, proportional sharing, real-time**2 Fast detection of communication patterns in distributed executions**

Thomas Kunz, Michiel F. H. Seuren

November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research****Publisher:** IBM PressFull text available: [pdf\(4.21 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

3 Implicit coscheduling: coordinated scheduling with implicit information in distributed systems

Andrea Carol Arpacı-Dusseau

August 2001 **ACM Transactions on Computer Systems (TOCS)**, Volume 19 Issue 3**Publisher:** ACM PressFull text available: [pdf\(1.83 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In modern distributed systems, coordinated time-sharing is required for communicating processes to leverage the performance of switch-based networks and low-overhead protocols. Coordinated time-sharing has traditionally been achieved with gang scheduling or explicit coscheduling, implementations of which often suffer from many deficiencies:

multiple points of failure, high context-switch overheads, and poor interaction with client-server, interactive, and I/O -intensive workloads. I ...

Keywords: clusters, coscheduling, gang scheduling, networks of workstations, proportional-share scheduling, two-phase waiting

4 An ultra low-power processor for sensor networks

 Virantha Ekanayake, Clinton Kelly, Rajit Manohar

October 2004 **ACM SIGOPS Operating Systems Review , ACM SIGARCH Computer Architecture News , ACM SIGPLAN Notices , Proceedings of the 11th international conference on Architectural support for programming languages and operating systems ASPLOS-XI**, Volume 38 , 32 , 39 Issue 5 , 5 , 11

Publisher: ACM Press

Full text available:  pdf(437.23 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present a novel processor architecture designed specifically for use in low-power wireless sensor-network nodes. Our sensor network asynchronous processor (SNAP/LE) is based on an asynchronous data-driven 16-bit RISC core with an extremely low-power idle state, and a wakeup response latency on the order of tens of nanoseconds. The processor instruction set is optimized for sensor-network applications, with support for event scheduling, pseudo-random number generation, bitfield operations, and ...

Keywords: asynchronous, event-driven, low-energy, picojoule computing, sensor network processor, sensor networks, wireless

5 Scheduling computations on a software-based router

 Xiaohu Qie, Andy Bavier, Larry Peterson, Scott Karlin

June 2001 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2001 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '01**, Volume 29 Issue 1

Publisher: ACM Press

Full text available:  pdf(1.46 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Recent efforts to add new services to the Internet have increased the interest in software-based routers that are easy to extend and evolve. This paper describes our experiences implementing a software-based router, with a particular focus on the main difficulty we encountered: how to schedule the router's CPU cycles. The scheduling decision is complicated by the desire to differentiate the level of service for different packet flows, which leads to two fundamental conflicts: (1) assigning proce ...

6 The design, implementation and evaluation of SMART: a scheduler for multimedia applications

 Jason Nieh, Monica S. Lam

October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5

Publisher: ACM Press

Full text available:  pdf(2.48 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

7 Efficient descriptor-vector multiplications in stochastic automata networks

 Paulo Fernandes, Brigitte Plateau, William J. Stewart

May 1998 **Journal of the ACM (JACM)**, Volume 45 Issue 3

Publisher: ACM Press

Full text available:  pdf(275.57 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper examines numerical issues in computing solutions to networks of stochastic

automata. It is well-known that when the matrices that represent the automata contain only constant values, the cost of performing the operation basic to all iterative solution methods, that of matrix-vector multiply, is given by $rN=i=1$

N

$n \times i = 1$

N

ni

Keywords: Markov chains, generalized tensor algebra, stochastic automata networks, vector-descriptor multiplication

8 [Borrowed-virtual-time \(BVT\) scheduling: supporting latency-sensitive threads in a general-purpose scheduler](#)

Kenneth J. Duda, David R. Cheriton

December 1999 **ACM SIGOPS Operating Systems Review , Proceedings of the seventeenth ACM symposium on Operating systems principles SOSP '99**, Volume 33 Issue 5

Publisher: ACM Press

Full text available:  [pdf\(1.81 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Systems need to run a larger and more diverse set of applications, from real-time to interactive to batch, on uniprocessor and multiprocessor platforms. However, most schedulers either do not address latency requirements or are specialized to complex real-time paradigms, limiting their applicability to general-purpose systems. In this paper, we present *Borrowed-Virtual-Time (BVT) Scheduling*, showing that it provides low-latency for real-time and interactive applications yet weighted sharin ...

9 [An object-based programming model for shared data](#)

Gail E. Kaiser, Brent Hailpern

April 1992 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 14 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(3.28 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#), [review](#)

The classical object model supports private data within objects and clean interfaces between objects, and by definition does not permit sharing of data among arbitrary objects. This is a problem for real-world applications, such as advanced financial services and integrated network management, where the same data logically belong to multiple objects and may be distributed over multiple nodes on the network. Rather than give up the advantages of encapsulated objects in modeling real-world en ...

Keywords: coordination language, daemons, financial applications, object-based, real-time, sharing

10 [Experience Using Multiprocessor Systems—A Status Report](#)

Anita K. Jones, Peter Schwarz

June 1980 **ACM Computing Surveys (CSUR)**, Volume 12 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(4.48 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

11 [Session 6: threads: Power-aware resource allocation in high-end systems via online simulation](#)

Barry Lawson, Evgenia Smirni

**June 2005 Proceedings of the 19th annual international conference on
Supercomputing ICS '05**

Publisher: ACM Press

Full text available:  pdf(298.15 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

Traditionally, scheduling in high-end parallel systems focuses on how to minimize the average job waiting time and on how to maximize the overall system utilization. Despite the development of scheduling strategies that aim at maximizing system utilization, parallel supercomputing traces that span long time periods indicate that such systems are mostly underutilized. Much of the time there is simply not enough load to keep the system fully utilized, although time periods do exist where system ut ...

Keywords: online simulation, parallel workload characterization, performance evaluation, power-aware scheduling, resource allocation

12 Improving Gang Scheduling through job performance analysis and malleability 

 Julita Corbalan, Xavier Martorell, Jesus Labarta

June 2001 **Proceedings of the 15th international conference on Supercomputing**

Publisher: ACM Press

Full text available:  pdf(150.61 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The OpenMP programming model provides parallel applications a very important feature: job malleability. Job malleability is the capacity of an application to dynamically adapt its parallelism to the number of processors allocated to it. We believe that job malleability provides to applications the flexibility that a system needs to achieve its maximum performance. We also defend that a system has to take its decisions not only based on user requirements but also based on run-time performance ...

13 Sensor databases: The design of an acquisitional query processor for sensor 

 networks

Samuel Madden, Michael J. Franklin, Joseph M. Hellerstein, Wei Hong

June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available:  pdf(485.52 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We discuss the design of an acquisitional query processor for data collection in sensor networks. Acquisitional issues are those that pertain to where, when, and how often data is physically acquired (*sampled*) and delivered to query processing operators. By focusing on the locations and costs of acquiring data, we are able to significantly reduce power consumption over traditional passive systems that assume the *a priori* existence of data. We discuss simple extensions to SQL for co ...

14 A survey of commercial parallel processors 

 Edward Gehringer, Janne Abullarade, Michael H. Gulyen

September 1988 **ACM SIGARCH Computer Architecture News**, Volume 16 Issue 4

Publisher: ACM Press

Full text available:  pdf(2.96 MB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

This paper compares eight commercial parallel processors along several dimensions. The processors include four shared-bus multiprocessors (the Encore Multimax, the Sequent Balance system, the Alliant FX series, and the ELXSI System 6400) and four network multiprocessors (the BBN Butterfly, the NCUBE, the Intel iPSC/2, and the FPS T Series). The paper contrasts the computers from the standpoint of interconnection structures, memory configurations, and interprocessor communication. Also, the share ...

15 Progress-based regulation of low-importance processes 

John R. Douceur, William J. Bolosky

December 1999 **ACM SIGOPS Operating Systems Review , Proceedings of the**

**seventeenth ACM symposium on Operating systems principles SOSP
'99, Volume 33 Issue 5**

Publisher: ACM Press

Full text available:  [pdf\(1.53 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

MS Manners is a mechanism that employs progress-based regulation to prevent resource contention with low-importance processes from degrading the performance of high-importance processes. The mechanism assumes that resource contention that degrades the performance of a high-importance process will also retard the progress of the low-importance process. MS Manners detects this contention by monitoring the progress of the low-importance process and inferring resource contention from a drop in the p ...

Keywords: process priority, progress-based feedback, symmetric resource contention

16 System-level power optimization: techniques and tools

 Luca Benini, Giovanni de Micheli

April 2000 ACM Transactions on Design Automation of Electronic Systems (TODAES),

Volume 5 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(385.22 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This tutorial surveys design methods for energy-efficient system-level design. We consider electronic systems consisting of a hardware platform and software layers. We consider the three major constituents of hardware that consume energy, namely computation, communication, and storage units, and we review methods of reducing their energy consumption. We also study models for analyzing the energy cost of software, and methods for energy-efficient software design and compilation. This survey ...

17 Disk-directed I/O for MIMD multiprocessors

 David Kotz

February 1997 ACM Transactions on Computer Systems (TOCS), Volume 15 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(559.18 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Many scientific applications that run on today's multiprocessors, such as weather forecasting and seismic analysis, are bottlenecked by their file-I/O needs. Even if the multiprocessor is configured with sufficient I/O hardware, the file system software often fails to provide the available bandwidth to the application. Although libraries and enhanced file system interfaces can make a significant improvement, we believe that fundamental changes are needed in the file server software. We prop ...

Keywords: MIMD, collective I/O, disk-directed I/O, file caching, parallel I/O, parallel file system

18 DHT: OpenDHT: a public DHT service and its uses

 Sean Rhea, Brighten Godfrey, Brad Karp, John Kubiatowicz, Sylvia Ratnasamy, Scott Shenker, Ion Stoica, Harlan Yu

August 2005 Proceedings of the 2005 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '05

Publisher: ACM Press

Full text available:  [pdf\(535.74 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Large-scale distributed systems are hard to deploy, and distributed hash tables (DHTs) are no exception. To lower the barriers facing DHT-based applications, we have created a public DHT service called OpenDHT. Designing a DHT that can be widely shared, both among mutually untrusting clients and among a variety of applications, poses two distinct challenges. First, there must be adequate control over storage allocation so that greedy or

malicious clients do not use more than their fair share. Se ...

Keywords: distributed hash table, peer-to-peer, resource allocation

19 CROP: cluster resource optimization package for PVM applications



Everett E. Mullis

April 1998 **Proceedings of the 36th annual Southeast regional conference**

Publisher: ACM Press

Full text available: [pdf\(1.34 MB\)](#) Additional Information: [full citation](#), [references](#), [index terms](#)

20 Energy aware design: Managing static leakage energy in microprocessor functional units



Steven Dropsho, Volkan Kursun, David H. Albonesi, Sandhya Dwarkadas, Eby G. Friedman
November 2002 **Proceedings of the 35th annual ACM/IEEE international symposium on Microarchitecture**

Publisher: IEEE Computer Society Press

Full text available: [pdf\(1.21 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)
[Publisher Site](#)

Static energy due to subthreshold leakage current is projected to become a major component of the total energy in high performance microprocessors. Many studies so far have examined and proposed techniques to reduce leakage in on-chip storage structures. In this study, static energy is reduced in the integer functional units by leveraging the unique qualities of dual threshold voltage domino logic. Domino logic has desirable properties that greatly reduce leakage current while providing fast prop ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)